

METHOD FOR AVOIDING DELAYS DURING SNOOP REQUESTS

CROSS REFERENCE TO CO-PENDING APPLICATIONS

A1
5 The present application is related to co-pending U.S. Patent
Application Serial No. 09/626,030, filed July 27, 2000, entitled
Cooperative Hardware and Microcode Control System for Pipelined
Instruction Execution; U.S. Patent Application Serial No.
08/650,800, filed August 30, 2000, entitled Method for Improved
10 First Level Cache Coherency; U.S. Patent Application Serial No.
09/650,730, filed August 30 2000, entitled Leaky Cache Mechanism;
and U.S. Patent Application Serial No. 08/235,196, filed April 29,
1994, entitled Data Coherency Protocol for Multi-level Cached High
Performance Multiprocessor System, assigned to the assignee of the
15 present invention and incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention: - The present invention relates
generally to data processing systems employing multiple instruction
20 processors and more particularly relates to multiprocessor data
processing systems employing multiple levels of cache memory.

2. Description of the Prior Art: - It is known in the art
that the use of multiple instruction processors operating out of

FIG. 3 is a more detailed block diagram of Voyager IP 50, located within Subpod 28, located within POD 20 (see also Figs. 1 and 2). As explained above, each instruction processor has a dedicated system controller having a dedicated level two cache memory. Instruction processor 64 has two dedicated level one cache memories (not shown in this view). One level one cache memory is a read-only memory for program instruction storage. Instruction processor 64 executes its instructions from this level one cache memory. The other level one cache memory (also not shown in this view) is a read/write memory for operand storage.

Instruction processor 64 is coupled via its two level one cache memories and dedicated system controller 58 to the remainder of the system. System controller 58 contains input logic 74 to interface with instruction processor 64. In addition, data path logic 68 controls movement of the data through system controller 58. The utilitarian functions are provided by Locks, Dayclocks, and UPI 62.

The remaining elements of system controller 58 provide the level two cache memory functions. SLC data ram 66 is the data actual storage facility. Control logic 70 provides the cache management function. SLC tags 72 are the tags associated with the level two cache memory. FLC-IC Dup. Tags 76 provides the duplicate tags for the level one instruction cache memory of instruction

FIG. 3 is a more detailed block diagram of Voyager IP 50, located within Subpod 28, located within POD 20 (see also Figs. 1 and 2). As explained above, each instruction processor has a dedicated system controller having a dedicated level two cache memory. Instruction processor 64 has two dedicated level one cache memories (not shown in this view). One level one cache memory is a read-only memory for program instruction storage. Instruction processor 64 executes its instructions from this level one cache memory. The other level one cache memory (also not shown in this view) is a read/write memory for operand storage. ***

Instruction processor 64 is coupled via its two level one cache memories and dedicated system controller 58 to the remainder of the system. System controller 58 contains input logic 74 to interface with instruction processor 64. In addition, data path logic 68 controls movement of the data through system controller 58. The utilitarian functions are provided by Locks, Dayclocks, and UPI 62.

The remaining elements of system controller 58 provide the level two cache memory functions. SLC data ram 66 is the data actual storage facility. Control logic 70 provides the cache management function. SLC tags 72 are the tags associated with the level two cache memory. FLC-IC Dup. Tags 76 provides the duplicate tags for the level one instruction cache memory of instruction

Fig. 5 is a timing diagram of an example of the operation of the present invention. In the preferred embodiment, the instruction processor timing clock cycles are divided into sub-cycles, called slots. In the present view, time proceeds from left to right in the timing diagram. Each vertical column represents one slot. The two horizontal rows associated with Tag slots 94 and Data slots 96 represent the prior art solution in response to data and SNOOP requests, whereas the two horizontal rows associated with Tag slots 98 and Data slots 100 represent the savings associated with the present invention.

In the prior art solution, a first data request is made (i.e., Req. 1) which occupies both the tag slot and data slot for the first time slot. Similarly, the second request occupies both the tag slot and data slot for the second time slot. The request 1 and request two updates occupy the third and fourth time slots. This happens even though a SNOOP request is present. Thus, the SNOOP request is acted upon only during the ninth tag time slot, because there is no new data request present. This delay of the SNOOP request means that other processors within the system remain idle waiting for the response to the SNOOP request, even though the local processor has all of its requests honored on a high priority basis.

According to the present invention, however, as shown in tag